21.2 Clock Generation and Distribution of a Dual-Core Xeon® Processor with 16MB L3 Cache

Simon Tam, Justin Leung, Rahul Limaye, Sam Choy, Sujal Vora, Mitsuhiro Adachi

Intel, Santa Clara, CA

The processor described in this paper is a dual-core Xeon® processor fabricated on a 65nm CMOS process with 8 levels of copper interconnect and 16MB of integrated on-die L3 cache (LLC) [1]. The processor embodies 1.328 billion transistors in a die area of 435mm^2. In this paper, the clock generation and clock distribution of this processor are described.

Figure 21.2.1 shows the physical organization and the clock domain architecture of the processor. The processor contains CORE0 and CORE1. The un-core encompasses the LLC, the I/O pads, and other logic blocks. The critical clocking features of this processor are: (1) multiple clock domains with different frequencies and (2) dedicated core and un-core voltage domains.

The two cores of this processor operate synchronously at the MCLK frequency which is an integer multiple of the 200MHz system clock (BCLK = MCLK/N, where N is an integer). The processor cores support a wide range of MCLK multiplication factors and operate at more than 3.0GHz at a 1.25V core supply voltage (Vcore). The L3 cache and the associated un-core logic operate at $^1/_2$ the core frequency (SCLK = $^1/_2$MCLK) using a separate un-core voltage supply (Vcache). Two identical interfaces integrated with the un-core logic and operating at the SCLK frequency are dedicated for the core to the un-core data flow. An independent ZCLK clock domain operating at 4 times the system clock frequency serves the front-side-bus (FSB). At the 200MHz system clock frequency, the FSB of this processor is capable of 800MT/s.

To support the two cores and the two clock domains in the un-core while maintaining good phase alignment between the core and un-core clock, this processor employs a parallel clock generator topology in which the differential system clock (BCLK) serves as the common reference. Figure 21.2.1 shows the physical topology (the PLL label represents the clock generator). Close proximity of these generators enables a package-level daisy-chain-style routing of the BCLK to the corresponding reference clock inputs. On-die reference clock de-skew fuses are used to compensate for the BCLK delay offsets that may appear at the reference clock inputs. With fuse compensation, the design achieves less than 40ps of BCLK mismatch.

Figure 21.2.2 shows the un-core clock generator architecture and the un-core clock distributions. The core and un-core use identical clock generators but have different customized clock distribution networks. The clock generator encompasses two PLLs: the IOPLL and the CorePLL. They are responsible for the generation of the ZCLK and the MCLK, respectively [2]. The IOPLL in this cascaded architecture isolates the CorePLL from BCLK noise to attain a quieter core clock. In the core and the un-core, an S-MACRO library cell samples the MCLK rising edges to generate the half frequency SCLK. The placements of the S-MACRO library cells are inside the local circuit blocks in the core and, for simplicity, are embedded within the pre-global clock distribution in the un-core. The core clock distribution is un-changed from the original single-core design [3]. It employs a recombinant MCLK pre-global clock tree implemented with CMOS inverters driving a MCLK global clock grid. Inverters used in the core pre-global distribution help to maintain the duty-cycle fidelity.

The un-core operates at $^1/_2$ the MCLK frequency and minimizing the un-core clock power and distribution delay are the critical design objectives. As a result, a hybrid clock distribution consisting of a MCLK pre-global network and a SCLK global grid is chosen for the un-core distribution architecture. The un-core pre-global clock tree consists of two horizontal spines and nine vertical spines (Fig. 21.2.3). The clock spines contain both the pre-global and global clock drivers. Figure 21.2.4 shows the complete un-core clock distribution hierarchy. The un-core clock distribution exhibits 900ps of delay and consumes about 6.7W of clock-distribution power. The un-core hybrid clock distribution represents a 33% reduction in distribution delay and a 3X reduction in clock power when compared to a core-style-based reference design. Figure 21.2.7 shows the relative delay profile of the un-core global clock. The simulated skew is less than 11ps.

The un-core clock distribution incorporates fuse-based de-skew buffers to gain additional skew control [4]. De-skew buffers are strategically located at the root of each of the vertical spines. The de-skew buffer is implemented using a CMOS delay chain, selectable by either fuses or TAP control. The de-skew buffer exhibits 60ps of range as measured, selectable in 4 steps.

Data transfer between the core and the un-core interface will need to cross both the clock domain and the V_{CC} domain simultaneously. The parallel clock generator architecture in this processor results in an implicit alignment of the core and the un-core clocks. To minimize the impacts of inter-clock-domain uncertainties caused by clock generator and voltage mismatches, a pipelined de-skew logic (PDSL) circuit with embedded voltage level shifter (VLS) is dedicated for the core to un-core data flow. The PDSL circuit (Fig. 21.2.5) consists of three flip-flop latch pairs integrated with the VLS circuit. The PDSL interface circuit is designed to handle high core to un-core clock uncertainties without failure. Each direction of the core to un-core data flow is served by one set of PDSLs.

Within the un-core, data from the I/O pads operating at the ZCLK frequency and the un-core logic operating at the SCLK frequency must cross the second major inter-clock domain transfer pipe. The ZCLK and the SCLK are related to each other at the ratio of 8 to N. The ZCLK-to-SCLK domain-crossing interface needs to accommodate this ratio without degrading FSB data rate while having good immunity to inter-clock domain uncertainties. Figure 21.2.6 shows the circuit structure deployed to support this ZCLK-to-SCLK domain-crossing. A similar circuit exists that is dedicated for the SCLK-to-ZCLK crossing. This interface employs parallel flip-flops to transmit and receive the data at the full FSB data rate. In this ZCLK-to-SCLK example, specific enables are generated for the transmit flip-flop and the receive flip-flop. Figure 21.2.6 shows the corresponding clock and data waveforms

Acknowledgements:
The authors acknowledge Matt Cross, Jeff McMillan, Jaspriya Baweja, Beatriz Jimenez, and Grace Sun for their contributions.

References:
[1] S.Rusu, et al., "A Dual-Core Multi-Threaded Xeon® Processor with 16MB L3 Cache," *ISSCC Dig. Tech. Papers,* Paper 5.3, Feb., 2006.
[2] K.Wong, et al., "Cascaded PLL Design for a 90nm CMOS High Performance Microprocessor," *ISSCC Dig. Tech. Papers,* pp. 422-423, Feb., 2003.
[3] N.Bindal, et al., "Scalable Sub-10ps Skew Global Clock Distribution for a 90nm Multi-GHz IA Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 346-347, Feb., 2003.
[4] S.Tam, et al., "Clock generation and Distribution for the Third Generation Itanium® Processor," *Symp. VLSI Circuits*, pp. 9-12, 2003.

 1-4244-0079-1/06/$20.00

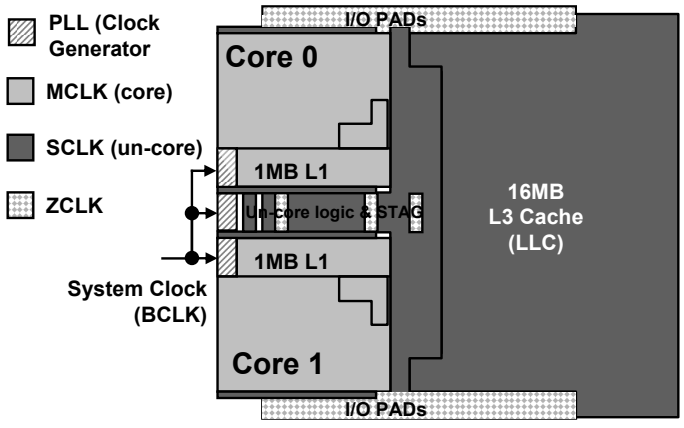

Figure 21.2.1: Clock domain and physical organization.

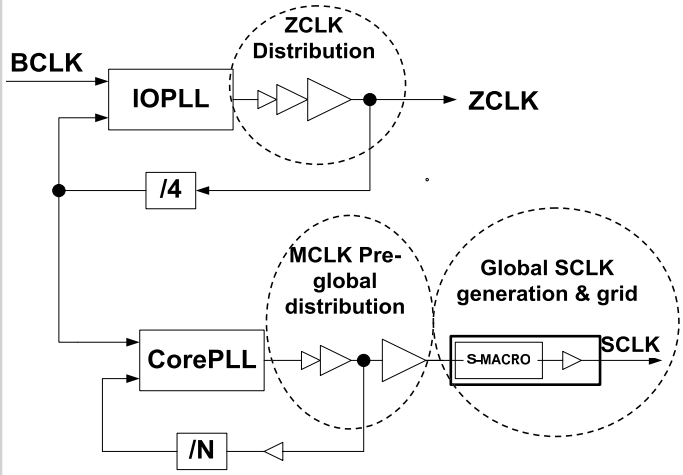

Figure 21.2.2: Clock-generator architecture.

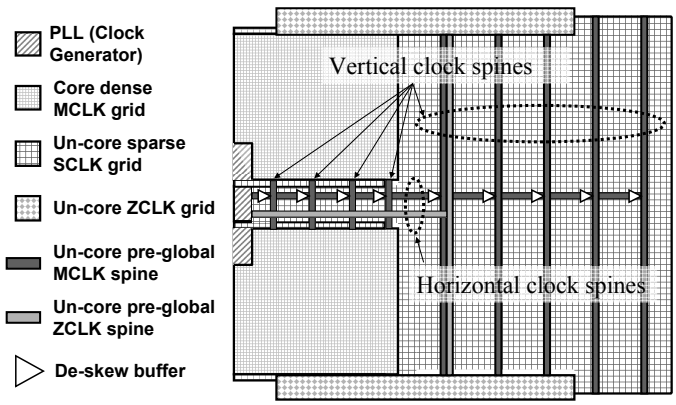

Figure 21.2.3: Clock distribution.

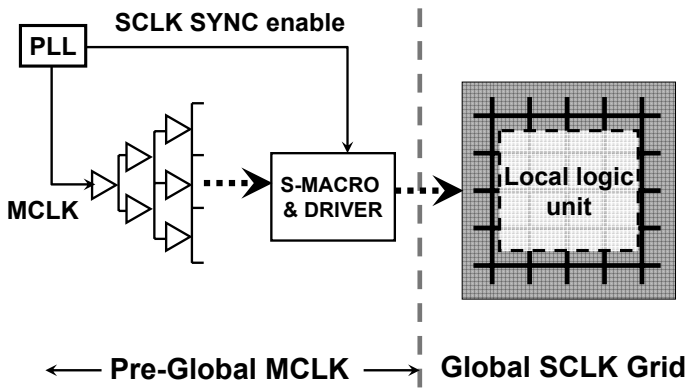

Figure 21.2.4: Clock hierarchy.

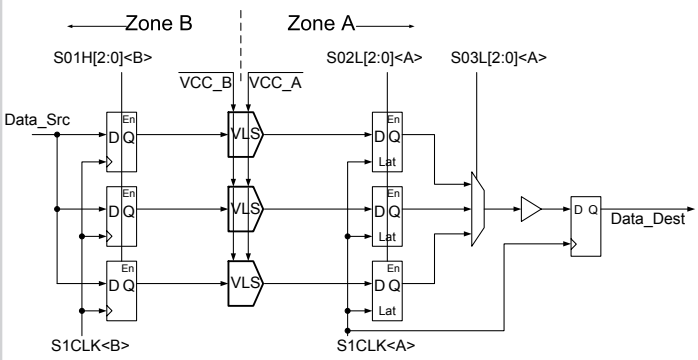

Figure 21.2.5: Pipelined de-skew logic at the core to un-core interface.

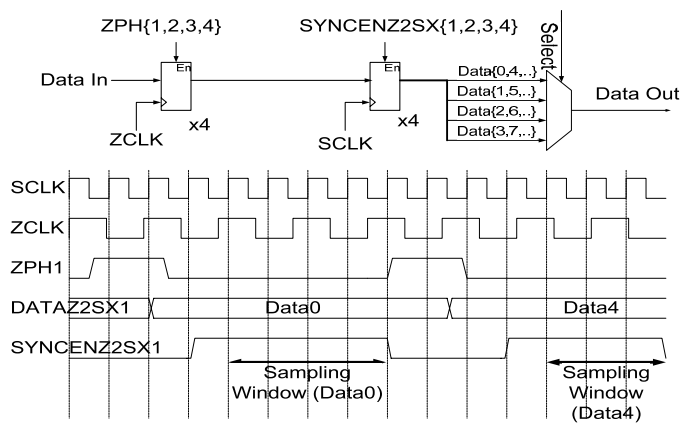

Figure 21.2.6: Front-side-bus to un-core clock domain crossing interface.

Continued on Page 660

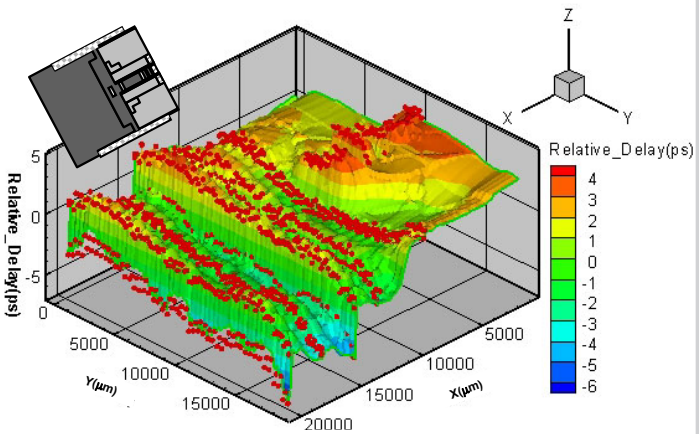

Figure 21.2.7: Relative clock delay of the n-core global clock.

